arXiv:1510.02961v1 [math.OC] 10 Oct 2015

Sparse plus Low rank Network Identification: A Nonparametric Approach ⋆

Mattia Zorzi [a], Alessandro Chiuso [a]

[a] *Dipartimento di Ingegneria dell'Informazione, Università degli studi di Padova, via Gradenigo 6/B, 35131 Padova, Italy*

Abstract

Modeling and identification of high-dimensional stochastic processes is ubiquitous in many fields. In particular, there is a growing interest in modeling stochastic processes with simple and interpretable structures. In many applications, such as econometrics and biomedical sciences, it seems natural to describe each component of that stochastic process in terms of few latent variables, which are not accessible for observation, and possibly of few other components of the stochastic process. These relations can be encoded in graphical way via a structured dynamic network, referred to as "sparse plus low-rank (S+L) network" hereafter. The problem of finding the S+L network as well as the dynamic model can be posed as a system identification problem. In this paper, we introduce two new nonparametric methods to identify dynamic models for stochastic processes described by a S+L network. These methods take inspiration from regularized estimators based on recently introduced kernels (e.g. "stable spline", "tuned-correlated" etc.). Numerical examples show the benefit to introduce the S+L structure in the identification procedure.

Key words: Linear system identification, Sparsity and low-rank inducing priors, Kernel-based methods, Gaussian processes.

1 Introduction

In many applications, high-dimensional data are measured to describe the underlying phenomena. These data are typically measured over time and can thus be modeled as a stochastic process whose components are referred to as *manifest* variables, i.e. accessible for observation. It is often reasonable to assume that these observations are somehow linked to an (unmeasurable and small) set of variables, called *latent* variables. The aim of these latent variables is to capture common relations (comovements) to all the manifest variables. Relations among manifest variables and latent variables can be encoded through a two-layer graph (or network) where the nodes in the top (and hidden) layer denote the latent variables, whereas the ones in the bottom (and visible) layer denote the manifest variables. The direct relation, in a sense to be precisely defined later on, between two nodes is encoded by the presence of a connecting directed edge. If the relations among manifest variables are mostly encoded through the latent variables and the number of the latter is small (as compared to the number of manifest variables), this graph will be referred to as *sparse plus low-rank (S+L) network*; in fact, as we shall see, its structure translates into a S+L structure for the dynamic model of the manifest process. In particular, the rank of the low-rank component coincides with the number of latent variables whereas the sparse component depends on the number of edges among the nodes representing the manifest variables. This modeling framework seems natural in many applications, such as econometrics (Chandrasekaran et al., 2010; Choi et al., 2010; Zorzi & Sepulchre, 2015), biomedical sciences (Liegegois et al., 2015), and so on.

In the first part of this paper we propose a new S+L network defined as a directed graph wherein edges encode conditional Granger causality dependences (Granger, 1969; Ding et al., 2006) among variables. In the special case where there is no latent variable our model coincides with the sparse model considered in Chiuso & Pillonetto (2012). On the contrary, when the relations among manifest variables are only characterized through the latent variables we obtain a quasi-static factor model where the noise process is white (Heij et al., 1997; Deistler & Zinner, 2007; Deistler et al., 2015). Therefore, our S+L model can be understood as blend of those two models.

Within this framework, we shall formulate an identification problem to select, using a finite set of measured

⋆ This work has been partially supported by the FIRB project "Learning meets time" (RBFR12M3AC) funded by MIUR.

Email addresses: zorzimat@dei.unipd.it (Mattia Zorzi), chiuso@dei.unipd.it (Alessandro Chiuso).

data, the most appropriate S+L model (and thus the corresponding network).

A consolidated paradigm in system identification is the so-called prediction error method (PEM), see Ljung (1999); Söderström & Stoica (1988). In the traditional setting, candidate models are described in fixed parametric model structures, e.g. ARMAX. However, there are two main difficulties in this setting. First a model selection problem (i.e. order selection), usually performed by AIC and BIC criteria (Akaike, 1974; Schwarz, 1978), needs to be solved. Second, the parametrization of the predictor is non-linear, so that minimizing the squared prediction error leads to a non-convex optimization problem. Regularization has been recently introduced in the PEM framework, see Pillonetto et al. (2011); Pillonetto & De Nicolao (2010); Chen et al. (2012); Chiuso & Pillonetto (2012); Pillonetto et al. (2014), as an alternative approach to control complexity of the estimated models. With this latter setting we search the candidate model, described via the predictor impulse responses, in an infinite dimensional nonparametric model class; the inverse (ill-posed) problem of determining a specific model using a finite set of measured data can be made into a well posed one using a penalty term, whose duty is to favor models with specific features. In the Bayesian view, this is equivalent to the introduction of an *a priori* probability (prior) on the unknown model. In the nonparametric Gaussian regression approach proposed in Pillonetto et al. (2011), this prior distribution is completely characterized by the covariance function, known also as kernel in the machine learning literature. The kernel encodes the *a priori* knowledge about the predictor impulse responses. In our case, the *a priori* knowledge is that the predictor impulse responses must be Bounded Input Bounded Output (BIBO) stable and respect the S+L structure. Starting from these *a priori* assumptions, we derive the corresponding kernel by using the maximum entropy principle. In particular, we consider two possible alternative formulations to endow this *a priori* properties in the kernel, and thus two different identification algorithms.

These kernels are characterized by the decay rate of the predictor impulse responses, by the number of conditional Granger causality relations among the manifest variables and by the number of latent variables. This ensemble of features is not known and is characterized by the so called hyperparameters vector. The latter is usually estimated by minimizing the negative log-likelihood of the measured data (Rasmussen & Williams, 2006). In our case, the challenge is to perform the joint estimation of the hyperparameters tuning the sparse part and the low-rank one. Indeed, it should be observed that the S+L decomposition of a given model might not be unique. On the other hand, once the hyperparameters vector is fixed the uniqueness of the estimated model will be guaranteed through regularization. To estimate the hyperparameters minimizing the negative log-likelihood, we propose an algorithm imposing an "hyper regularizer" on the low-rank hyperparameter to partially handle the non-uniqueness of the S+L decomposition.

Numerical experiments involving both S+L and generic models show the effectiveness of our identification procedure both in terms of complexity, predictive capability and impulse response fit.

The outline of the paper follows. In Section 2, we introduce the S+L models. In Section 3, we introduce the S+L identification problem. In Section 4, we derive the maximum entropy kernels inducing BIBO stability, sparsity and low-rank. In Section 5, we present our S+L identification procedure. In Section 6, we provide some numerical examples to show the effectiveness of our method. Finally, conclusions are drawn in Section 7. In order to streamline the presentation all proofs are deferred to the Appendix.

Notation

Throughout the paper, we will use the following notation. $\mathbb{N}^+$ is the set of natural numbers without zero. $\mathbb{E}[\cdot]$ denotes the expectation, while $\mathbb{E}[\cdot|\cdot]$ denotes the best linear estimator, i.e. conditional expectation in the Gaussian case. Given three (possibly infinite dimensional) random vectors a, b and c we say that a is conditionally uncorrelated of b given c if

$$\mathbb{E}[\mathrm{a}|\mathrm{b},\mathrm{c}] = \mathbb{E}[\mathrm{a}|\mathrm{c}].$$

Given $G \in \mathbb{R}^{n\times p}$, $[G]_{ij}$ denotes the entry of G in position (i,j). $\mathcal{M}^p$ denotes the vectors space of symmetric matrices of dimension p. $\mathcal{M}^p_+$ is the cone of symmetric positive definite matrices of dimension p, and $\overline{\mathcal{M}}^p_+$ denotes its closure. $\ell_2(\mathbb{N}^+)$ denotes the space of $\mathbb{R}$-valued infinite length sequences, which we think as infinite dimensional column vectors $g := [g_1\ g_2\ \dots\ g_j\ \dots]^\top$, $g_k \in \mathbb{R}$, $k \in \mathbb{N}^+$, such that $\|g\|_2 := \sqrt{\sum_{k=1}^{\infty} |g_k|^2} < \infty$. $\ell_2^{p\times n}(\mathbb{N}^+)$ is the space of semi-infinite matrices

$$\Phi = \begin{bmatrix} (\phi^{[11]})^\top & \dots & (\phi^{[1n]})^\top \\ \vdots & \ddots & \vdots \\ (\phi^{[p1]})^\top & \dots & (\phi^{[pn]})^\top \end{bmatrix} \tag{1}$$

where $\phi^{[ij]} \in \ell_2(\mathbb{N}^+)$, $i = 1\dots p$ and $j = 1\dots n$. $\ell_1(\mathbb{N}^+)$ denotes the space of $\mathbb{R}$-valued infinite length sequences g such that $\|g\|_1 := \sum_{k=1}^{\infty} |g_k| < \infty$. $\ell_1^{p\times n}(\mathbb{N}^+)$ is defined in similar way. $\mathcal{S}_2(\mathbb{N}^+)$ denotes the space of symmetric infinite dimensional matrices K such that $\|K\|_2 := \sqrt{\sum_{i,j=1}^{\infty} |[K]_{ij}|^2} < \infty$. $\mathcal{S}_2^p(\mathbb{N}^+)$ is the space of symmet-

ric infinite dimensional matrices

$$K = \begin{bmatrix} K^{[11]} & K^{[12]} & \dots & K^{[1p]} \\ K^{[12]} & K^{[22]} & & K^{[2p]} \\ \vdots & & \ddots & \vdots \\ K^{[1p]} & K^{[2p]} & \dots & K^{[pp]} \end{bmatrix} \tag{2}$$

where $K^{[ij]} \in \mathcal{S}_2(\mathbb{N}^+)$, $i,j = 1 \dots p$. Given $\Phi \in \ell_2^{p\times n}(\mathbb{N}^+)$, $\Psi \in \ell_2^{m\times n}(\mathbb{N}^+)$ and $K \in \mathcal{S}_2^n(\mathbb{N}^+)$, the products $\Phi\Psi^\top$ and $\Phi K \Psi^\top$ are understood as $p \times m$ matrices whose entries are limits of infinite sequences (Jorgesen et al., 2011). Given a transfer matrix $L(z)$, $z \in \mathbb{C}$, of dimension $p \times p$, with some abuse of terminology, we say that $L(z)$ has rank n, with $n \leq p$, if it admits the decomposition $L(z) = FH(z)$ where F is a $p \times n$ matrix and $H(z)$ is a $n \times p$ transfer matrix. Given a stochastic process $y = \{y(t)\}_{t\in\mathbb{Z}}$, its i-th component is denoted by $y_i = \{y_i(t)\}_{t\in\mathbb{Z}}$. With some abuse of notation, $y(t)$ will both denote a random vector and its sample value. From now on the time t will denote *present* and we shall talk about *past* and *future* with respect to time t. With this convention in mind,

$$\mathrm{y}^- = \begin{bmatrix} y(t-1)^\top & y(t-2)^\top & \dots \end{bmatrix}^\top \tag{3}$$

denotes the (infinite length) past data vector of y at time t. In similar way, y_i^- denotes the past data vector of y_i at time t.

2 Sparse plus Low rank Models

Consider two zero-mean jointly stationary and Gaussian stochastic processes y and x taking values in $\mathbb{R}^p$ and $\mathbb{R}^n$, respectively. Let y be manifest, i.e. it can be measured, and x latent, i.e. it is not accessible for observation. In this paper we shall assume that y and x are jointly described by the model

$$\begin{aligned} y(t) =& Fx(t) + S(z)y(t) + v(t) \\ x(t) =& H(z)y(t) + w(t) \end{aligned} \tag{4}$$

where $H(z) = \sum_{k=1}^\infty H_k z^{-k}$ is a BIBO stable transfer matrix of dimension $n \times p$, $F \in \mathbb{R}^{p\times n}$, $S(z) = \sum_{k=1}^\infty S_k z^{-k}$ is a BIBO stable transfer matrix of dimension $p \times p$; v and w are, respectively, p and n dimensional, zero mean and independent white Gaussian noises (WGN) with covariance matrices Σ_v and Σ_w.

It is possible to describe the structure of model (4) using a directed graph (Lauritzen, 1996) having two layers. The nodes in the top layer correspond to the scalar pocesses x_i $i = 1 \dots n$, i.e. the components of x, whereas the ones in the bottom layer correspond to the scalar processes y_i $i = 1 \dots p$, i.e. the components of y, see Figure 1. Then, the connections among the nodes obey the following rules:

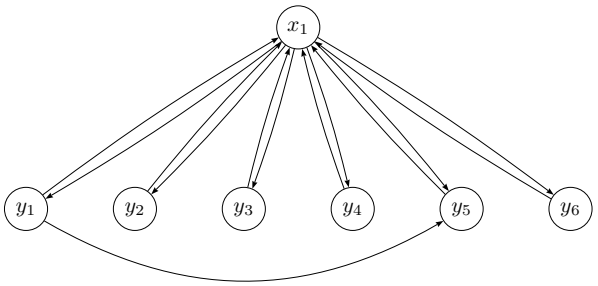

Fig. 1. Example of a network describing a S+L model with $p = 6$ manifest variables and $n = 1$ latent variable.

- there is a directed link from node y_j to node y_i if

  $$\mathbb{E}[y_i(t)|x(t), \mathrm{y}^-] \neq \mathbb{E}[y_i(t)|x(t), \mathrm{y}_k^-, k = 1, .., p, k \neq j]$$

  i.e. if y_j^- is needed to predict $y_i(t)$ given y_k^-, $\forall k \neq j$ and $x(t)$. In this case, we shall say y_j conditionally Granger causes y_i
- there is a directed link from node y_j to node x_i if

  $$\mathbb{E}[x_i(t)|\mathrm{y}^-] \neq \mathbb{E}[x_i(t)|\mathrm{y}_k^-, k = 1, .., p, k \neq j]$$

  i.e. if y_j^- is needed to predict $x_i(t)$ given y_k^- with $k \neq j$. In this case, we shall say y_j conditionally Granger causes x_i
- there is a direct link from node x_j to node y_i if

  $$\mathbb{E}[y_i(t)|\mathrm{y}^-, x(t)] \neq \mathbb{E}[y_i(t)|\mathrm{y}^-, x_k(t), k = 1, .., n, k \neq j]$$

  i.e. if $x_j(t)$ is needed to predict $y_i(t)$ given y^- and $x_k(t)$ with $k \neq j$. In this case, we shall say x_j conditionally Granger causes y_i.

Our main modeling assumption in (4) is that manifest variables Granger cause each other mostly through few latent variables. Therefore, we have $n \ll p$, i.e. the number of latent variables is small as compared to the number of manifest variables, and $S(z)$ is sparse, i.e. there are few directed edges connecting the manifest variables $y_j(t)$. From (4), we obtain the sparse plus low-rank (S+L) model for y:

$$y(t) = S(z)y(t) + L(z)y(t) + e(t) \tag{5}$$

where $S(z)$ is a sparse transfer matrix by assumption, $L(z) := FH(z)$ is a low-rank transfer matrix since F is tall ($n \ll p$), and $e(t) := v(t) + w(t)$ is WGN with covariance matrix $\Sigma = \Sigma_v + F\Sigma_w F^\top$.

Let $\hat{y}(t|t-1)$ be the minimum variance one-step ahead predictor of $y(t)$ based on the past data y^-, and $\hat{x}(t|t-1)$

be the minimum variance estimator of $x(t)$ based on y^-. From (4), we have

$$\begin{aligned}\hat{y}(t|t-1) &= F\hat{x}(t|t-1) + S(z)y(t)\\ \hat{x}(t|t-1) &= H(z)y(t)\end{aligned} \tag{6}$$

that is, the predictable part of y is a function of the (estimated) latent variables and of the pasts of few manifest variables. Of course equations (6) can be compacted so that $\hat{y}(t|t-1)$ can be expressed explicitly as a function of its own past:

$$\hat{y}(t|t-1) = S(z)y(t) + L(z)y(t) \tag{7}$$

via a S+L model. In Figure 1 we provide an example of a network describing a S+L model with $p = 6$ manifest variables and $n = 1$ latent variable. In particular, y_5 is conditionally Granger caused by x and y_1, y_i with $i \neq 5$ is conditionally Granger caused only by x, and x is conditionally Granger caused by y. Therefore, the corresponding S+L model in (5) has $S(z)$ with only one nonnull entry in position (5,1) and $L(z)$ has rank equal to one.

In the case that $L(z) = 0$, i.e. there is no need of latent variables to characterize the predictor of y, we obtain the sparse model presented in Chiuso & Pillonetto (2012). In the case $S(z) = 0$, i.e. the predictor of y is completely characterized by the estimators of the latent variables. More precisely, we obtain a quasi-static factor model where F is the factor loading matrix, x is the n-dimensional factor and e is WGN (Deistler & Zinner, 2007). Accordingly, the S+L model (4) is a blend of the two above models. It is worth noting that process x can be understood as a state in the special case that $S(z) = 0$. In the general case (i.e. $S(z) \neq 0$) x is a pseudo state and the aim is to place in $S(z)$ few (scalar) transfer functions whose state dimension is very large. In this way x has dimension a lot smaller than the state process. We conclude that model (4) is a generalization of the state space model.

The decomposition of a transfer matrix into sparse plus low-rank may not be unique. As noticed in Chandrasekaran et al. (2010), this degeneracy may occur when $L(z)$ is sparse, i.e. the latent variables are not sufficiently "diffuse" across the manifest variables, or the degree of sparsity of $S(z)$ is low, i.e. there are manifest variables conditionally Granger caused by too many manifest variables. Let $\mathcal{V}_S(s)$ be the manifold of transfer matrices having at most s nonnull entries, and $\mathcal{V}_L(n)$ the manifold of transfer matrices having rank less than or equal to n. Given $G(z) \in \mathcal{V}_S(s) + \mathcal{V}_L(n)$, its S+L decomposition is typically non unique, that is its sparse part in $\mathcal{V}_S(s)$ and its low-rank part in $\mathcal{V}_L(n)$ are not identifiable. On the other hand, if $G(z) = S(z) + L(z)$ with $S(z) \in \mathcal{V}_S(s)$, $L(z) \in \mathcal{V}_L(n)$, and the tangent spaces of $\mathcal{V}_S(s)$ and $\mathcal{V}_L(n)$, respectively, at $S(z)$ and $L(z)$ have transverse intersection then the S+L decomposition of $G(z)$ is locally identifiable, (Chandrasekaran et al., 2010, Section 3). We do not tackle this nonidentifiability issue in this paper, although it is important. Indeed, our aim is to find one S+L decomposition (see the next Section) which is not necessarily unique.

3 Problem Formulation

Assume measured data $\{y(t)\}_{t=1,..,N}$ are available from the manifest process y generated by (4). The latent process x cannot be measured nor its dimension n is known. In this section, we address the problem of estimating $S(z)$ and $L(z)$ in (5) from $\{y(t)\}_{t=1,..,N}$.

The transfer matrix $S(z)$ is parametrized in terms of its impulse response coefficients S_k, which are stacked in $\theta_s^\top \in \ell_1^{1\times p^2}(\mathbb{N}^+)$ defined as follows:

$$\begin{aligned}\theta_s^\top = \Big[&(s^{[11]})^\top \ \dots \ (s^{[1p]})^\top \Big| \dots\\ &\dots \Big| (s^{[p1]})^\top \ \dots \ (s^{[pp]})^\top \Big]\end{aligned} \tag{8}$$

where $s^{[ij]} \in \ell_1(\mathbb{N}^+)$ is such that

$$s^{[ij]} = \Big[\ [S_1]_{ij} \ [S_2]_{ij} \ \dots \ [S_k]_{ij} \ \dots \Big]^\top. \tag{9}$$

In a similar way, the impulse response coefficients L_k parametrizing the transfer matrix $L(z)$ are stacked in $\theta_l^\top \in \ell_1^{1\times p^2}(\mathbb{N}^+)$ as follows

$$\begin{aligned}\theta_l^\top = \Big[&(l^{[11]})^\top \ \dots \ (l^{[1p]})^\top \Big| \dots\\ &\dots \Big| (l^{[p1]})^\top \ \dots \ (l^{[pp]})^\top \Big]\end{aligned} \tag{10}$$

where $l^{[ij]} \in \ell_1(\mathbb{N}^+)$ is such that

$$l^{[ij]} = \Big[\ [L_1]_{ij} \ [L_2]_{ij} \ \dots \ [L_k]_{ij} \dots \Big]^\top. \tag{11}$$

The measured data $y(1) \dots y(N)$ are stacked in the vector y^+ as follows

$$\begin{aligned}\mathrm{y}^+ = \Big[&y_1(t)^\top \ \dots \ y_1(t+N-1)^\top \Big| \dots\\ &\dots \Big| y_p(t)^\top \ \dots \ y_p(t+N-1)^\top \Big]^\top\end{aligned} \tag{12}$$

and similarly

$$\begin{aligned}\mathrm{e}^+ = \Big[&e_1(t)^\top \ \dots \ e_1(t+N-1)^\top \Big| \dots\\ &\dots \Big| e_p(t)^\top \ \dots \ e_p(t+N-1)^\top \Big]^\top.\end{aligned} \tag{13}$$

Let $\Phi \in \ell_2^{pN \times p^2}(\mathbb{N}^+)$ be the regression matrix defined as

$$\Phi = I_p \otimes \begin{bmatrix} \phi_1 & \dots & \phi_p \end{bmatrix} \tag{14}$$

where $\phi_j \in \ell_2^{N \times 1}(\mathbb{N}^+)$ is the *Toeplitz* matrix defined as

$$[\phi_j]_{kh} = y_j(t + k - h - 1) \tag{15}$$

with $k = 1 \dots N$ and $h \in \mathbb{N}^+$. From (5) the vector y^+ containing the measured data satisfy the linear model [1]

$$\mathrm{y}^+ = \underbrace{\Phi(\theta_l + \theta_s)}_{\hat{\mathrm{y}}^+} + \mathrm{e}^+ \tag{16}$$

where $\hat{\mathrm{y}}^+ := \Phi(\theta_s + \theta_l)$ is the one-step ahead predictor of y^+.

Therefore, our S+L identification problem can be formulated in terms of PEM as follows.

Problem 1 *Find $\theta_s^\top, \theta_l^\top \in \ell_1^{1 \times p^2}(\mathbb{N}^+)$ corresponding to a S+L model and minimizing the prediction error norm $\|\mathrm{y}^+ - \Phi(\theta_s + \theta_l)\|^2_{\Sigma^{-1} \otimes I_N}$.*

Following the nonparametric Gaussian regression approach in Pillonetto et al. (2011), we model θ_s and θ_l as two zero-mean processes with kernel $K_S \in \mathcal{S}_2^{p^2}(\mathbb{N}^+)$ and $K_L \in \mathcal{S}_2^{p^2}(\mathbb{N}^+)$, respectively. Moreover, we shall assume that θ_s and θ_l are Gaussian and independent. We will see in Section 4 the latter assumptions are suggested by the maximum entropy principle. In this Gaussian framework, the optimal solution to Problem 1 is the conditional mean, given the measured data, of $S(z)$ and $L(z)$, see Section 5. Before proceeding, the kernels K_S and K_L have to be designed in such a way to enforce the posterior mean $\hat{S}(z)$ of $S(z)$ to be BIBO stable and sparse, and the posterior mean $\hat{L}(z)$ of $L(z)$ to be BIBO stable and low-rank.

4 Maximum entropy kernels

In this Section we characterize the probability law of the joint process $[\theta_s^\top \ \theta_l^\top]^\top$, satisfying the desired properties, by using the maximum entropy principle. Our unique assumption on that process is that it is zero-mean and admitting a joint probability density. We will see that, under suitable constraints, the optimal solution (i.e. maximizing the differential entropy) is a Gaussian process where θ_s and θ_l are independent. Then, we will also characterize the corresponding kernels. In what follows, we propose two different ways to enforce BIBO stability and low-rank on $\hat{L}(z)$. This leads to two different types of kernel for θ_l.

4.1 First type of kernel

We start with the constraints on θ_s inducing BIBO stability and sparsity on $\hat{S}(z)$. We consider the constraint on θ_s

$$\mathbb{E}[\|s^{[ij]}\|^2_{P^{-1}}] \le c_{ij}, \tag{17}$$

where $c_{ij} \ge 0$, $i, j = 1 \dots p$, and $P \in \mathcal{S}_2(\mathbb{N}^+)$. If $c_{ij} = 0$, then $s^{[ij]}$ is the null sequence in mean square and so is its posterior mean.

Proposition 2 *Let ϕ be a zero-mean process with kernel $K \in \mathcal{S}_2(\mathbb{N}^+)$. Consider the constraint*

$$\mathbb{E}[\|\phi\|^2_{P^{-1}}] \le c \tag{18}$$

where $c \ge 0$, and $P \in \mathcal{S}_2(\mathbb{N}^+)$ strictly positive definite kernel (in the sense of Moore) such that $[P]_{tt} \le \kappa t^\alpha e^{-\beta t}$, $t \in \mathbb{N}^+$, with $\kappa, \beta > 0$ and $\alpha \in \mathbb{R}$. Then constraint (18) implies that for any $\varepsilon > 1$ there exists $\bar{\kappa}_\varepsilon > 0$ such that

$$[K]_{tt} \le \bar{\kappa}_\varepsilon t^{\alpha+\varepsilon} e^{-\beta t}, \ \ t \in \mathbb{N}^+. \tag{19}$$

It is not difficult to see that the assumption on P in Proposition 2 is satisfied by the kernels usually employed in the identification of dynamical models (e.g. stable spline, tuned/correlated and so on, see Pillonetto et al. (2014)). Thus, by (17) the covariance of the k-th element of $s^{[ij]}$ decays exponentially. We conclude that the posterior mean of the transfer function in position (i, j) of $S(z)$, under constraint (17), is BIBO stable. Moreover, it is null if and only if $c_{ij} = 0$.

Remark 3 *Clearly, if P in* (17) *is chosen as the covariance matrix of DC (TC) prior, the resulting Maximum Entropy prior coincides with the DC (TC) prior. Some recent literature has discussed the maximum entropy properties of some kernels, such as the TC or DC kernels, see Nicolao et al. (1998), Carli et al. (2014) and Chen et al. (2015), as well as extensions to more articulated kernels Prando et al. (2015). In these papers it was shown for TC and DC kernels that weaker constraints with respect to* (17) *lead to the same solution. We refer the reader to these papers form more details; here we stick to the simpler conditions* (17) *in order to streamline the derivation as well as to allow alternative choices of the matrix P in* (17).

[1] The matrix Φ is to be build from measured data and, as such, it is only partially available. In practice model (16) needs to be approximated truncating Φ (and thus θ_l and θ_s). This corresponds to assuming zero initial conditions, and it is a reasonable approximation given the decay, as a function of k, of the coefficients L_k and S_k (which is a necessary condition for BIBO stability). We shall not enter into these details in the paper. See also Remark 14 in Section 5.

In view of Proposition 2, to guarantee BIBO stability on $\hat{L}(z)$ we impose the constraint

$$\sum_{i,j=1}^{p} \mathbb{E}[\|l^{[ij]}\|^2_{P^{-1}}] \leq r \tag{20}$$

for some $r > 0$.

We now switch our attention to the low rank property of the matrix $L(z)$. Let A_l be the random semi-infinite matrix defined as

$$A_l = \begin{bmatrix} L_1 & L_2 & \dots & L_k & \dots \end{bmatrix} \tag{21}$$

and consider the constraint on θ_l

$$\mathbb{E}[A_l A_l^\top] \leq Q \tag{22}$$

with $Q \in \overline{\mathcal{M}}^p_+$. If Q has $p-n$ singular values equal to zero, so does the covariance $\mathbb{E}[A_l A_l^\top]$; therefore the posterior mean $\hat{A}_l$ of A_l has rank equal to n thus admitting the decomposition

$$\hat{A}_l = \begin{bmatrix} \hat{F}\hat{H}_1 & \hat{F}\hat{H}_2 & \dots & \hat{F}\hat{H}_k & \dots \end{bmatrix}, \tag{23}$$

where $\hat{F} \in \mathbb{R}^{p\times n}$ and the $\hat{H}_k \in \mathbb{R}^{n\times p}$, $k \in \mathbb{N}^+$, as in Section 2. We conclude that, under constraints (20), $\hat{L}(z)$ is BIBO stable and under constraint (22), its rank is less than or equal to n if and only if Q has rank equal to n.

In order to build the desired prior distribution we make use of Kolmogorov extension Theorem, see Øksendal (1998), and work with finite vectors extracted from processes θ_l and θ_s.

To do so, let us consider a finite index set $\mathcal{I} = \mathcal{I}_s \times \mathcal{I}_l$ in $\mathbb{N}^+ \times \mathbb{N}^+$. Let $\check{\theta}_s$ be the random vector whose components are extracted from the process θ_s according to the index set $\mathcal{I}_s$. In similar way, we define the random vector $\check{\theta}_l$ whose components are extracted from θ_l according to $\mathcal{I}_l$. We denote by $\mathbf{p}_\mathcal{I}(\check{\theta}_s, \check{\theta}_l)$ the joint probability density of $\check{\theta}_s$ and $\check{\theta}_l$. By Kolmogorov extension Theorem the joint process $[\theta_s^\top \ \theta_l^\top]^\top$ can be characterized by specifying the joint probability density $\mathbf{p}_\mathcal{I}$ for all finite sets $\mathcal{I} \subset \mathbb{N}^+ \times \mathbb{N}^+$. Accordingly, our idea is to characterize the joint process through $\mathbf{p}_\mathcal{I}$ using the maximum entropy principle. Such principle states that among all the probability densities $\mathbf{p}_\mathcal{I}$ satisfying the desired constraints, the optimal one should maximize the differential entropy (Cover & Thomas, 1991)

$$\mathbf{H}(\mathbf{p}_\mathcal{I}) = -\mathbb{E}[\log(\mathbf{p}_\mathcal{I})]. \tag{24}$$

Constraints (17), (20) and (22) boil down, respectively, to

$$\mathbb{E}[\|\check{s}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_s}}] \leq c_{ij}, \quad i,j = 1\dots p \tag{25}$$

$$\sum_{i,j=1}^{p} \mathbb{E}[\|\check{l}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_l}}] \leq r \tag{26}$$

$$\mathbb{E}[\check{A}_l \check{A}_l^\top] \leq Q \tag{27}$$

where $\check{s}^{[ij]}$ and $\check{l}^{[ij]}$ are the vectors extracted from $s^{[ij]}$ and $l^{[ij]}$ according to the index set $\mathcal{I}_s$ and $\mathcal{I}_l$, respectively. $P_{\mathcal{I}_s}$ and $P_{\mathcal{I}_l}$ are the kernel matrices whose entries are extracted from P according to $\mathcal{I}_s$ and $\mathcal{I}_l$, respectively. $\check{A}_l$ is the matrix whose blocks are extracted from A_l according to $\mathcal{I}_l$. Therefore, we obtain the following maximum entropy problem

$$\begin{aligned} &\max_{\mathbf{p}_\mathcal{I} \in \mathcal{P}} \mathbf{H}(\mathbf{p}_\mathcal{I}) \\ &\text{s.t. } \mathbb{E}[\|\check{s}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_s}}] \leq c_{ij}, \quad i,j = 1\dots p \\ &\qquad \sum_{i,j=1}^{p} \mathbb{E}[\|\check{l}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_l}}] \leq r \\ &\qquad \mathbb{E}[\check{A}_l \check{A}_l^\top] \leq Q \end{aligned} \tag{28}$$

where $\mathcal{P}$ is the space of probability densities in $\mathbb{R}^{|\mathcal{I}_s|\times|\mathcal{I}_l|}$ which are Lebesgue integrable.

Theorem 4 *Under the assumption that $c_{ij} > 0$, $i,j = 1\dots p$, and $Q \in \mathcal{M}^p_+$, the unique optimal solution to the maximum entropy problem (28) is such that $\check{\theta}_s$ and $\check{\theta}_l$ are independent, Gaussian with zero mean and kernel matrix, respectively,*

$$\check{K}_S = \operatorname{diag}(\gamma_1 \dots \gamma_{p^2}) \otimes P_{\mathcal{I}_s} \tag{29}$$

$$\check{K}_L = (\lambda I_{p^2} \otimes P^{-1}_{\mathcal{I}_l} + \tilde{\Lambda} \otimes I_{p|\mathcal{I}_l|})^{-1} \tag{30}$$

where $\gamma_i > 0$, $i = 1\dots p^2$, $\tilde{\Lambda} \in \overline{\mathcal{M}}^p_+$ and $\lambda \geq 0$.

In what follows we assume that constraints (26) and (27) are totally binding in problem (28). Since λ and $\tilde{\Lambda}$ are the Lagrange multipliers associated to those constraints, it is not difficult to see that in this situation $\lambda > 0$ and $\tilde{\Lambda} \in \mathcal{M}^p_+$. Moreover, we define $\Lambda = \tilde{\Lambda}^{-1}$. As noticed before, we are interested in the limiting cases where c_{ij} might be equal to zero and Q might be of low-rank in order to have sparse and low-rank estimators. To include these scenarios, we consider the limits as $c_{ij} \to 0$ and Q tends to a low rank matrix and extend the maximum entropy solution by continuity.

Proposition 5 *Let $\mathbf{C} = \{(i,j) \text{ s.t. } c_{ij} = 0\}$ and $\mathbf{Q} = \{v \text{ s.t. } Qv = 0\}$. Then, the maximum entropy solution extended by continuity is the probability density such that*

$\check{\theta}_s$ and $\check{\theta}_l$ are independent, Gaussian, zero-mean, and with kernel matrices

$$\check{K}_S = \mathrm{diag}(\gamma_1 \dots \gamma_{p^2}) \otimes P_{\mathcal{I}_s} \tag{31}$$

$$\check{K}_L = \lambda^{-1} I_{p^2} \otimes P_{\mathcal{I}_l} - \lambda^{-2} I_{p^2} \otimes P_{\mathcal{I}_l} \times (\lambda^{-1} I_{p^2} \otimes P_{\mathcal{I}_l} + \Lambda \otimes I_{p|\mathcal{I}_l|})^{-1} I_{p^2} \otimes P_{\mathcal{I}_l} \tag{32}$$

where $\gamma_{(i-1)p+j} = 0$ if and only if $(i,j) \in \mathbf{S}$ and $\Lambda v = 0$ if and only if $v \in \mathbf{Q}$.

Finally, in view of Kolmogorov extension Theorem, from the probability density of $[\check{\theta}_s^\top \ \check{\theta}_l^\top]^\top$ we can characterize the probability law of $[\theta_s^\top \ \theta_l^\top]^\top$ maximizing the differential entropy.

Corollary 6 *Consider the joint process $[\theta_s^\top \ \theta_l^\top]^\top$ where θ_s and θ_l are Gaussian independent processes with kernel, respectively,*

$$K_S = \mathrm{diag}(\gamma_1 \dots \gamma_{p^2}) \otimes P$$

$$K_L = \lambda^{-1} I_{p^2} \otimes P - \lambda^{-2} I_{p^2} \otimes P \times (\lambda^{-1} I_{p^2} \otimes P + \Lambda \otimes I_\infty)^{-1} I_{p^2} \otimes P \tag{33}$$

where I_∞ such that $[I_\infty]_{tt} = 1$, $t \in \mathbb{N}^+$, and zero otherwise. For all finite sets $\mathcal{I} \subset \mathbb{N}^+ \times \mathbb{N}^+$, its joint probability density is the extended solution to the maximum entropy problem (28).

It is worth noting, that the maximum entropy kernels are characterized by the hyperparameters γ_i, $i = 1 \dots p^2$, which control sparsity on $\hat{S}(z)$, Λ tuning the rank on $\hat{L}(z)$, while β (see Proposition 2) controls the decay rate (as a function of the time index $k \in \mathbb{N}^+$) of the estimators $\hat{L}_k$ and $\hat{S}_k$ and thus BIBO stability of $\hat{S}(z)$ and $\hat{L}(z)$, see Pillonetto & De Nicolao (2010). Finally, λ represents a trade-off between BIBO stability and low-rank on $\hat{L}(z)$.

Remark 7 *It is possible to derive the same structure for K_L by adopting the regularization point of view. In Wipf (2012), it has been shown that the penalty term $\log\det(R)$, with $R \in \mathcal{M}_+^p$, induces low-rank on R. Moreover, the term $\log\det(R)$ admits the variational upper bound*

$$\log\det(R) \le \mathrm{tr}(\Lambda^{-1} R) - \log\det(\Lambda^{-1}) - p \tag{34}$$

where $\Lambda \in \mathcal{M}_+^p$ and equality holds if and only if $R = \Lambda$. Consider the random vector $\check{\theta}_l$ extracted from θ_l according to $\mathcal{I}_l$. Thus, we can induce low-rank on $R = \check{A}_l \check{A}_l^\top$ by considering the penalty

$$\mathrm{tr}(\Lambda^{-1} \check{A}_l \check{A}_l^\top) - \log\det(\Lambda) - p \tag{35}$$

where Λ represents a rough estimate of $\check{A}_l \check{A}_l^\top$. The unique term depending on $\check{\theta}_l$ in (35) is $\mathrm{tr}(\Lambda^{-1} \check{A}_l \check{A}_l^\top) = \|\check{\theta}_l\|^2_{\Lambda^{-1} \otimes I_{p|\mathcal{I}_l|}}$ which is one part of the norm $\|\check{\theta}_l\|^2_{\check{K}_L^{-1}}$ with $\check{K}_L$ kernel matrix defined in (30). Thus, this penalty terms induces low-rank on $\check{A}_l$.

4.2 Second type of kernel

The BIBO stability and low-rank constraints on $\hat{L}(z)$ can be imposed by using only one constraint. Consider the random semi-infinite matrix A_l defined in (21) and the constraint

$$\mathbb{E}[A_l (P^{-1} \otimes I_p) A_l^\top] \le Q \tag{36}$$

with $Q \in \overline{\mathcal{M}}_+^p$. Similarly to the previous case, if the null space of Q has dimension $m-n$, then the posterior mean $\hat{A}_l$ of A_l has rank less than or equal to n; therefore also $\hat{L}(z)$ has rank less than or equal to n.

Proposition 8 *Assume that $P \in \mathcal{S}_2(\mathbb{N}^+)$ is strictly positive definite and such that $[P]_{tt} \le \kappa t^\alpha e^{-\beta t}$, $t \in \mathbb{N}^+$, with $\kappa, \beta > 0$ and $\alpha \in \mathbb{R}$. Then, under constraint (36), for any $\varepsilon > 1$ there exists $\bar{\kappa}_\varepsilon > 0$ such that $\mathbb{E}[|[L_t]_{ij}|^2] \le \bar{\kappa}_\varepsilon t^{\alpha+\varepsilon} e^{-\beta t}$, $i,j = 1 \dots p$, $t \in \mathbb{N}^+$.*

Therefore, constraint (36) guarantees that $\hat{L}(z)$ is BIBO stable and with rank less than or equal to n.

Similarly to the previous case, the joint process $[\theta_s^\top \ \theta_l^\top]^\top$ is characterized through $\mathbf{p}_\mathcal{I}$ using the maximum entropy principle. In particular, constraint (36) becomes

$$\mathbb{E}[\check{A}_l (P_{\mathcal{I}_l}^{-1} \otimes I_p) \check{A}_l^\top] \le Q. \tag{37}$$

The corresponding maximum entropy problem is

$$\max_{\mathbf{p}_\mathcal{I} \in \mathcal{P}} \mathbf{H}(\mathbf{p}_\mathcal{I}) \quad \text{s.t. } \mathbb{E}[\|\check{s}^{[ij]}\|^2_{P_{\mathcal{I}_s}^{-1}}] \le c_{ij}, \ \ i,j = 1 \dots p \quad \mathbb{E}[\check{A}_l (P_{\mathcal{I}_l}^{-1} \otimes I_p) \check{A}_l^\top] \le Q \tag{38}$$

Theorem 9 *Under the assumption that $c_{ij} > 0$, $i,j = 1 \dots p$, and $Q \in \mathcal{M}_+^p$, the unique optimal solution to the maximum entropy problem (38) is such that $\check{\theta}_s$ and $\check{\theta}_l$ are independent, Gaussian with zero mean and kernel matrix, respectively,*

$$\check{K}_S = \mathrm{diag}(\gamma_1 \dots \gamma_{p^2}) \otimes P_{\mathcal{I}_s} \tag{39}$$

$$\check{K}_L = \Lambda \otimes I_p \otimes P_{\mathcal{I}_l} \tag{40}$$

where $\gamma_i > 0$, $i = 1 \dots p^2$, and $\Lambda \in \mathcal{M}_+^p$.

Similarly to the previous case, we extend the maximum entropy solution by continuity to the case of interest.

Proposition 10 *Let* $\mathbf{C} = \{(i,j)\ s.t.\ c_{ij} = 0\}$ *and* $\mathbf{Q} = \{v\ s.t.\ Qv = 0\}$. *Then, the maximum entropy solution extended by continuity is the probability density such that* $\breve{\theta}_s$ *and* $\breve{\theta}_l$ *are independent, Gaussian, zero-mean, and with kernel matrices (39) and (40) , respectively, where* $\gamma_{(i-1)p+j} = 0$ *if and only if* $(i,j) \in \mathbf{S}$ *and* $\Lambda v = 0$ *if and only if* $v \in \mathbf{Q}$.

Finally, from the probability density of $[\breve{\theta}_s^\top \ \breve{\theta}_l^\top]^\top$ we characterize the probability law of the joint process $[\theta_s^\top \ \theta_l^\top]^\top$ using the Kolmogorov extension Theorem.

Corollary 11 *Consider the joint process* $[\theta_s^\top \ \theta_l^\top]^\top$ *where* θ_s *and* θ_l *are Gaussian independent processes with kernel, respectively,*

$$\begin{aligned} K_S &= \mathrm{diag}(\gamma_1 \dots \gamma_{p^2}) \otimes P \\ K_L &= \Lambda \otimes I_p \otimes P. \end{aligned} \tag{41}$$

For all finite sets $\mathcal{I} \subset \mathbb{N}^+ \times \mathbb{N}^+$, *its joint probability density is the extended solution to the maximum entropy problem (38).*

Note that, in respect to the first type of kernel, we do not need the extra hyperparameter λ.

Remark 12 *In the maximum entropy kernels we derived the matrix* P, *satisfying the assumptions of Proposition 2 and Proposition 8, needs to be fixed. In what follows, we consider the filtered kernel proposed in Pillonetto et al. (2011)*

$$P = FP_{SS}F^\top \tag{42}$$

where $P_{SS} \in \mathcal{S}_2(\mathbb{N}^+)$ *is the stable spline (SS) kernel and* F *is a Toeplitz matrix buil with the impulse response of a second order system, which is used to model high frequency oscillations which the predictor impulse response might have due to "high" frequency zeros of the noise spectrum. Note that, instead of SS one might choose other types of kernel, such as diagonal, tuned/correlated, diagonal/correlated and so on.*

5 Nonparametric Gaussian regression approach to S+L identification

In the following $\mathcal{H}_S$ and $\mathcal{H}_L$ denote the reproducing Hilbert spaces (Aronszajn, 1950) of deterministic functions on $\mathbb{N}^+$, associated with K_S and K_L, with norm denoted by $\|\cdot\|_{K_S^{-1}}$ and $\|\cdot\|_{K_L^{-1}}$, respectively. Let $\bar{\xi} := \{\tau, \xi\}$ denote the vector of all the hyperparameters characterizing the maximum entropy kernel of Section 4, where τ denotes the hyperparameters of P defined in (42), $\xi := \{\gamma_1 \dots \gamma_{p^2}, \Lambda, \lambda\}$ for the first type of kernel, and $\xi := \{\gamma_1 \dots \gamma_{p^2}, \Lambda\}$ for the second one. We assume that the past data y^- neither affects the *a priori* probability on θ_s and θ_l nor carries information on $\bar{\xi}$ and Σ (Poggio & Girosi, 1990), that is

$$\begin{aligned} &\mathbf{p}(\mathrm{y}^+, \theta_l, \theta_s, \mathrm{y}^- | \bar{\xi}, \Sigma) \\ &\quad = \mathbf{p}(\mathrm{y}^+ | \theta_l, \theta_s, \mathrm{y}^-, \bar{\xi}, \Sigma) \mathbf{p}(\theta_l, \theta_s | \mathrm{y}^-, \bar{\xi}, \Sigma) \mathbf{p}(\mathrm{y}^- | \bar{\xi}, \Sigma) \\ &\quad \approx \mathbf{p}(\mathrm{y}^+ | \theta_l, \theta_s, \mathrm{y}^-, \bar{\xi}, \Sigma) \mathbf{p}(\theta_l, \theta_s | \bar{\xi}, \Sigma) \mathbf{p}(\mathrm{y}^-). \end{aligned} \tag{43}$$

Let

$$\hat{\theta}_s = \mathbb{E}[\theta_s | \mathrm{y}^+, \bar{\xi}, \Sigma], \ \ \hat{\theta}_l = \mathbb{E}[\theta_l | \mathrm{y}^+, \bar{\xi}, \Sigma] \tag{44}$$

be the posterior mean of θ_s and θ_l, respectively, given y^+, $\bar{\xi}$ and Σ. Note that, $\hat{\theta}_s$ and $\hat{\theta}_l$ are equivalent to $\hat{S}(z)$ and $\hat{L}(z)$, respectively. Since both processes are Gaussian, $\hat{\theta}_s$ and $\hat{\theta}_l$ are also the minimum variance estimators of θ_s and θ_l, respectively. The next Proposition shows that almost surely $\hat{\theta}_s$ and $\hat{\theta}_l$ belong to the spaces $\mathcal{H}_S$ and $\mathcal{H}_L$, respectively, and are solution to a *Tikhonov*-type variational problem.

Proposition 13 *Under the assumption that* y *is a stationary finite variance process and under approximation (43), almost surely we have*

$$\begin{aligned} (\hat{\theta}_s, \hat{\theta}_l) = \underset{\substack{\theta_s \in \mathcal{H}_{K_S} \\ \theta_l \in \mathcal{H}_{K_L}}}{\arg\min} &\|\mathrm{y}^+ - \Phi(\theta_s + \theta_l)\|^2_{\Sigma^{-1} \otimes I_N} \\ &+ \|\theta_s\|^2_{K_S^{-1}} + \|\theta_l\|^2_{K_L^{-1}}. \end{aligned} \tag{45}$$

Moreover, almost surely we have

$$\hat{\theta}_s = K_S \Phi^\top c, \ \ \hat{\theta}_l = K_L \Phi^\top c \tag{46}$$

where

$$c = (\Phi(K_S + K_L)\Phi^\top + \Sigma \otimes I_N)^{-1} \mathrm{y}^+. \tag{47}$$

Remark 14 *The semi-infinite regression matrix* Φ *depends on both* y^+ *and* y^-. *Since* y^- *is never completely known, a solution to handle the initial conditions consists of setting its unknown components to zero. In this way, the introduced error goes to zero as* N *increases (Ljung, 1999, Section 3.2). Alternatively it would be possible to incorporate initial conditions in the estimation problem, e.g. modeling also the free response of the system. This is however outside the scope of the paper and is only practically relevant when very slow dynamics are present.*

In order to compute $\hat{\theta}_s$ and $\hat{\theta}_l$, estimates of the noise variance Σ and of the hyperparameters vector $\bar{\xi}$ are needed. In this paper Σ is estimated using a low-bias ARX-model as suggested in (Goodwin et al., 1992). To estimate the hyperparameters of P we consider the unstructured model for y

$$y(t) = G(t)y(t) + e(t) \tag{48}$$

where $G(z) = \sum_{k=1}^{\infty} G_k z^{-k}$ is BIBO stable and e WGN with covariance matrix Σ. Equation (48) can be written as a linear regression model

$$\mathrm{y}^+ = \Phi\theta + \mathrm{e}^+ \tag{49}$$

where $\theta^\top \in \ell^{1\times p^2}(\mathbb{N}^+)$ contains the coefficients of $G(z)$. Following the Gaussian regression approach to estimate θ in the linear model (49), we model θ as a zero-mean Gaussian process with kernel $K = I_{p^2} \otimes P$. Therefore, the hyperparameters of P are estimated by minimizing the negative log-likelihood of y^+ according to (49), see Pillonetto et al. (2011).

Then the hyperparameters vector ξ can be estimated minimizing the negative log-marginal likelihood ℓ of y^+ according to (16); under the assumptions of Proposition 13, we have

$$\begin{aligned}\ell(\mathrm{y}^+, \xi) \;&=\; \tfrac{1}{2}\log\det V + \tfrac{1}{2}(\mathrm{y}^+)^\top V^{-1}\mathrm{y}^+ + \text{ const. term}\\ V \;&=\; \Phi(K_S + K_L)\Phi^\top + \Sigma\otimes I_N\end{aligned} \tag{50}$$

where K_L and K_S are defined in (33) or in (41) and the matrix P has been fixed as above. Notice that, the minimization of (50) with respect to ξ is a nonconvex constrained optimization problem. Accordingly, only local minima can be computed. However, the real challenge is to perform the joint estimation of the γ_i's and Λ because the sparse and the low-rank part might be not identifiable from the measured data. As a partial remedy it is useful to reduce the *degrees of freedom* of Λ (see also Prando et al. (2015)), constraining its structure as follows:

$$\Lambda = \alpha(I - UU^\top) + U\mathrm{diag}(\beta_1 \dots \beta_n)U^\top. \tag{51}$$

The matrix $U \in \mathbb{R}^{p\times n}$ is built, in view also of Remark 7, using the first n singular vectors of an estimate $\hat{A}_l\hat{A}_l^\top$ of $A_lA_l^\top$. In this way, the constraints in Λ are decoupled along the "most reliable" n singular vectors of $\hat{A}_l\hat{A}_l^\top$ and their orthogonal complement. This is equivalent to fix n variables (given from the estimate $\hat{A}_l\hat{A}_l^\top$) as latent variables. Regarding the hyperparameters γ_i's, in Aravkin et al. (2014) it has been shown that the minimization of (50) automatically leads sparsity. Therefore, the minimization of (50) is performed with respect to $\tilde{\xi} := \{\gamma_1 \dots \gamma_p, \alpha, \beta_1 \dots \beta_n, \lambda\}$ for the first type of kernel and with respect to $\tilde{\xi} := \{\gamma_1 \dots \gamma_p, \alpha, \beta_1 \dots \beta_n\}$ for the second one while n and U are fixed. The complete procedure to estimate n, U and $\tilde{\xi}$ we propose is outlined in Algorithm 1. $n^{(k)}$, $U^{(k)}$, $\tilde{\xi}^{(k)}$, $\hat{A}_l^{(k)}$ and $\hat{L}^{(k)}(z)$ denote, respectively, n, U, $\tilde{\xi}$, $\hat{A}_l$ and $\hat{L}(z)$ at the k-th iteration of the algorithm. The likelihood function (50), when Λ is constrained to be of the form (51), is a function of U. This dependence is made explicit in the notation

$$\begin{aligned}\ell(\mathrm{y}^+, \tilde{\xi}, U) &= \frac{1}{2}\log\det V + \frac{1}{2}(\mathrm{y}^+)^\top V^{-1}\mathrm{y}^+ + \text{ const. term}\\ V &= \Phi(K_S + K_L)\Phi^\top + \Sigma\otimes I_N\\ \Lambda &= \alpha(I - UU^\top) + U\mathrm{diag}(\beta_1 \dots \beta_n)U^\top\end{aligned}$$

where K_S and K_L are defined in (33) for the first type of kernel and in (41) for the second one. Finally, to efficiently compute a local minimum of (50) in the algorithm we used the scaled gradient projection algorithm developed in Bonettini et al. (2015).

Algorithm 1 Computation of n, U and $\tilde{\xi}$

$k = 0$
$n^{(0)} \leftarrow 0$
$U_{OPT}^{(0)} \leftarrow$ empty matrix
$\tilde{\xi}^{(0)} \leftarrow \underset{\tilde{\xi}}{\mathrm{argmin}}\, \ell(\mathrm{y}^+, \tilde{\xi}, U_{OPT}^{(0)})$
$\tilde{\xi}_{OPT}^{(0)} \leftarrow \tilde{\xi}^{(0)}$
repeat
  $k \leftarrow k + 1$
  $n^{(k)} \leftarrow n^{(k-1)} + 1$
  if $n^{(k)} = 1$ **then**
    $\hat{A}_l^{(k)} \leftarrow [\,\hat{G}_1 \;\; \hat{G}_2 \;\; \dots\,]$ where $\hat{G}_1, \hat{G}_2, \dots$ are the coefficients of $\hat{G}(z)$ estimated from (48) with $K = I_{p^2} \otimes P$
  else
    $\hat{A}_l^{(k)} \leftarrow [\,\hat{L}_1^{(k)} \;\; \hat{L}_2^{(k)} \;\; \dots\,]$ where $\hat{L}_1^{(k)}, \hat{L}_2^{(k)}, \dots$ are the coefficients of $\hat{L}^{(k)}(z)$ estimated from (45) with K_S and K_L having hyperparameters given by $U_{OPT}^{(n^{(k)}-1)}$ and $\tilde{\xi}_{OPT}^{(n^{(k)}-1)}$
  end if
  $U^{(k)} \leftarrow$ first $n^{(k)}$ singular vectors of $\hat{A}_l^{(k)}\;\hat{A}_l^{(k)\,\top}$
  $\tilde{\xi}^{(k)} \leftarrow \underset{\tilde{\xi}}{\mathrm{argmin}}\, \ell(\mathrm{y}^+, \tilde{\xi}, U^{(k)})$
  repeat
    $\tilde{\xi}_{OPT}^{(n^{(k)})} \leftarrow \tilde{\xi}^{(k)}$
    $U_{OPT}^{(n^{(k)})} \leftarrow U^{(k)}$
    $k \leftarrow k + 1$
    $n^{(k)} \leftarrow n^{(k-1)}$
    $\hat{A}_l^{(k)} \leftarrow [\,\hat{L}_1^{(k)} \;\; \hat{L}_2^{(k)} \;\; \dots\,]$ where $\hat{L}_1^{(k)}, \hat{L}_2^{(k)}, \dots$ are the coefficients of $\hat{L}^{(k)}(z)$ estimated from (45) with K_S and K_L having hyperparameters given by $U^{(k-1)}$ and $\tilde{\xi}^{(k-1)}$
    $U^{(k)} \leftarrow$ first $n^{(k)}$ singular vectors of $\hat{A}_l^{(k)}\;\hat{A}_l^{(k)\,\top}$
    $\tilde{\xi}^{(k)} \leftarrow \underset{\tilde{\xi}}{\mathrm{argmin}}\, \ell(\mathrm{y}^+, \tilde{\xi}, U^{(k)})$
  until $\ell(\mathrm{y}^+, \tilde{\xi}^{(k)}, U^{(k)}) < \ell(\mathrm{y}^+, \tilde{\xi}^{(k-1)}, U^{(k-1)})$
until $\ell(\mathrm{y}^+, \tilde{\xi}_{OPT}^{(n^{(k)})}, U_{OPT}^{(n^{(k)})}) < \ell(\mathrm{y}^+, \tilde{\xi}_{OPT}^{(n^{(k)}-1)}, U_{OPT}^{(n^{(k)}-1)})$
$n \leftarrow n^{(k)} - 1$
$U \leftarrow U_{OPT}^{(n^{(k)}-1)}$
$\tilde{\xi} \leftarrow \tilde{\xi}_{OPT}^{(n^{(k)}-1)}$

It is worth noting that our algorithm is similar to the non-separable reweighting scheme proposed in Wipf & Nagarajan (2010) for solving a sparse Bayesian learning problem. That algorithm iteratively alternates the computation of the optimal estimate (in our case $\hat{L}(z)$) and the closed form update of the hyperparameters (in our case it is given by (51)).

6 Numerical experiments

We consider two Monte Carlo studies of 100 runs where at any run a manifest process y of dimension $m = 6$ is considered. For each run in the Monte Carlo experiments an identification data set of size 500 and a test set of size 1000 are generated. We compare the true model of y, denoted by TRUE, with the following estimators

- SL-I: this is the sparse plus low-rank estimator (46) with K_S and K_L as in (33)
- SL-II: this is the sparse plus low-rank estimator (46) with K_S and K_L as in (41)
- L-I: this is the low-rank estimator (46) with K_S set equal to zero and K_L as in (33)
- L-II: this is the low-rank estimator (46) with K_S set equal to zero and K_L as in (41)
- S: this is the sparse estimator (46) with K_L set equal to zero
- SS: this is the estimator based on model (48) where $G(z)$ is modeled as a zero-mean Gaussian process with kernel $K = I_{p^2} \otimes P$ and P defined in (42).

For implementation purposes, the impulse responses are truncated to a certain length T. The latter represents the "practical" length of those impulse responses (Pillonetto et al., 2011). Note that, the truncation does not introduce bias-variance tradeoff because the filtered kernel P forces the estimated impulse responses to decay exponentially. In these experiments we used $T = 50$.

The following performance indexes are considered:

- relative complexity of the estimated model. It is quantified with

$$C = \lim_{T\to\infty} \frac{\#\mathcal{E}_T}{m^2 T} \tag{52}$$

  where $\#\mathcal{E}_T$ is the minimum number of parameters needed to characterize the model with impulse responses truncated at T. For instance, for the S+L model (4) with $S(z)$ having s nonnull transfer functions and $L(z) = FH(z)$ having rank n, we have $\#\mathcal{E}_T = sT + mn(T + 1)$. Moreover, m^2T is the number of coefficients needed to characterize all the truncated impulse responses of model (48) without structure (i.e. neither sparsity nor low-rank is imposed). If we quantify the complexity of the S+L network as the number of edges among the manifest nodes and from the latent ones to the manifest ones, then it is not difficult to see that its complexity is equal to m^2C. Similar conclusions can be found for networks having only the sparse or the low-rank part. Therefore, the smaller C is, the simpler the network found is.
- one-step-ahead Coefficient of Determination, denoted by COD_1. Such index quantifies how much the test set variance is explained by the forecast, and is defined as:

$$\text{COD}_1 = 1 - \frac{\frac{1}{1000}\sum_{t=1}^{1000} \|y^{\text{test}}(t) - \hat{y}^{\text{test}}(t|t-1)\|^2}{\frac{1}{1000}\sum_{t=1}^{1000} \|y^{\text{test}}(t) - \bar{y}^{\text{test}}\|^2} \tag{53}$$

  where $\bar{y}^{\text{test}}$ denotes the sample mean of the test set data $y(1)^{\text{test}} \dots y^{\text{test}}(1000)$ and $\hat{y}^{\text{test}}(t|t-1)$ is the one-step ahead prediction computed using the estimated model. Notice that, the larger COD_1 is, the better predictive performance of the estimator is
- average impulse response fit

$$\text{AIRF} = 100\left(1 - \sqrt{\frac{\frac{1}{T}\sum_{k=1}^{T} \|G_k - \hat{G}_k\|^2}{\frac{1}{T}\sum_{k=1}^{T} \|G_k - \bar{G}\|^2}}\right) \tag{54}$$

  where G_k and $\hat{G}_k$ are the impulse response coefficients of the true and estimated model, respectively, and $\bar{G} = \frac{1}{T}\sum_{k=1}^{T} G_k$. In the case the true (estimated) model is S+L we have $G_k = S_k + L_k$ ($\hat{G}_k = \hat{S}_k + \hat{L}_k$).

In the first Monte Carlo experiment, y is generated through model (5). In particular, $L(z)$ is randomly generated with rank $n = 1$. $S(z)$ is randomly generated with 7 nonnull transfer functions. The position of those transfer functions is randomly chosen. The top panel of Figure 2 shows the relative complexity of the models obtained with the estimators SL-I, SL-II, L-I, L-II and S reporting the boxplots of the values of C after the 100 runs. One can see that L-II is the best estimator, then SL-I and S-II, then L-I and finally S. The top panel of Figure 3, reporting the boxplots of COD_1, shows the one-step predictive performance of the estimators is similar. The top panel of Figure 4 shows the boxplots of AIRF. It is clear that SL-I and SL-II outperform the others.

In the second experiment, y is generated through a generic (i.e. unstructured) model. The bottom panels of Figure 2, Figure 3 and Figure 4 show the performance of the employed estimators as before. One can see that the best estimators in terms of COD_1 are SL-I, SL-II and L-I. SL-I and SL-II provide simpler models than those obtained by L-I. On the other hand in terms of AIRF, SL-I and SL-II are slightly worse than L-I. These two Monte Carlo experiments suggest that the proposed S+L models and estimators (SL-I and SL-II) provide an effective way of estimating complex model, yielding

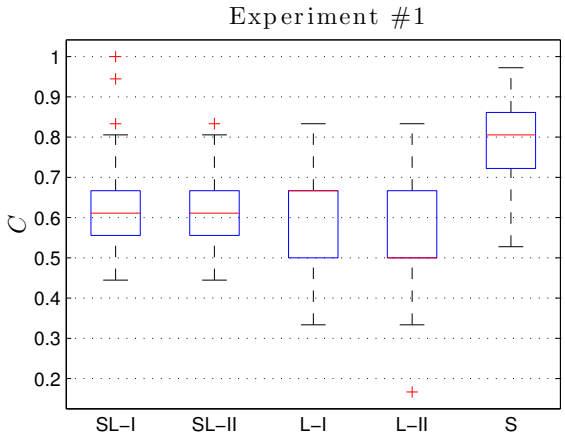

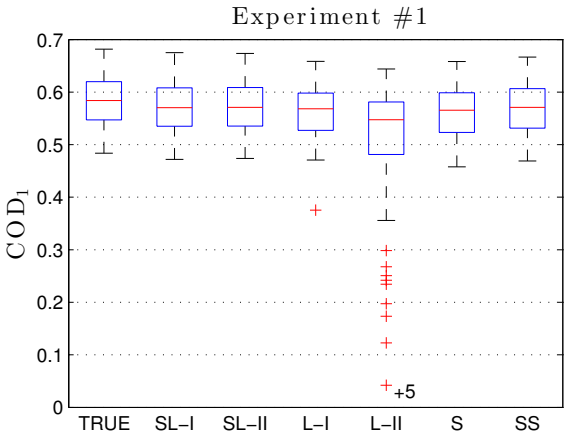

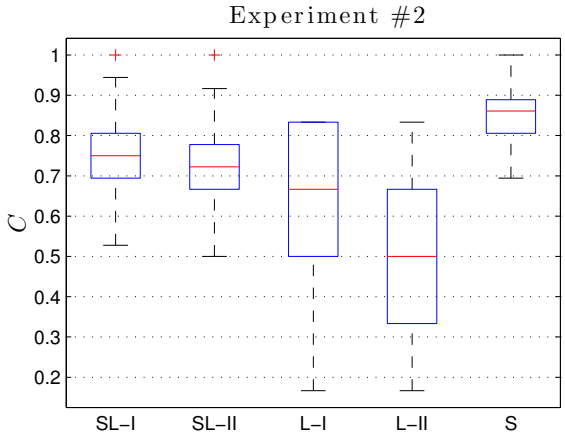

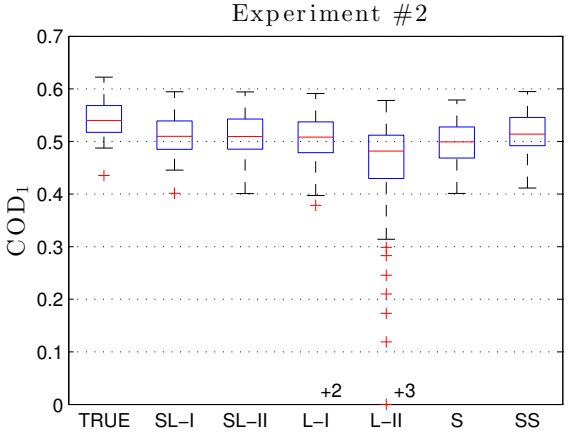

Fig. 2. Relative complexity of the models obtained by using the estimators SL-I, SL-II, L-I, L-II, and S.

Fig. 3. One step ahead coefficient of determination (COD_1) obtained by the 6 estimators described in Section 6. Moreover, TRUE provides an upper bound on the performance of those estimators, being the true model for the manifest process y.

a good tradeoff among model complexity, prediction accuracy and average impulse response fit.

7 Conclusions

In this paper, we proposed two procedures, based on a nonparametric Gaussian regression approach, to identify stochastic processes having a sparse plus low-rank (S+L) network. The kernels inducing the S+L structure have been derived using the maximum entropy principle. Simulations show that the proposed S+L estimators have good predictive capability as well as low complexity compared with the sparse estimator and the low-rank estimators.

Appendix

Proof of Proposition 2

Define $\breve{P} \in \mathcal{S}_2(\mathbb{N}^+)$ such that

$$[\breve{P}]_{ts} = \begin{cases} \kappa t^{\alpha} e^{-\beta t} & t = s \\ [P]_{ts} & \text{otherwise,} \end{cases} \tag{55}$$

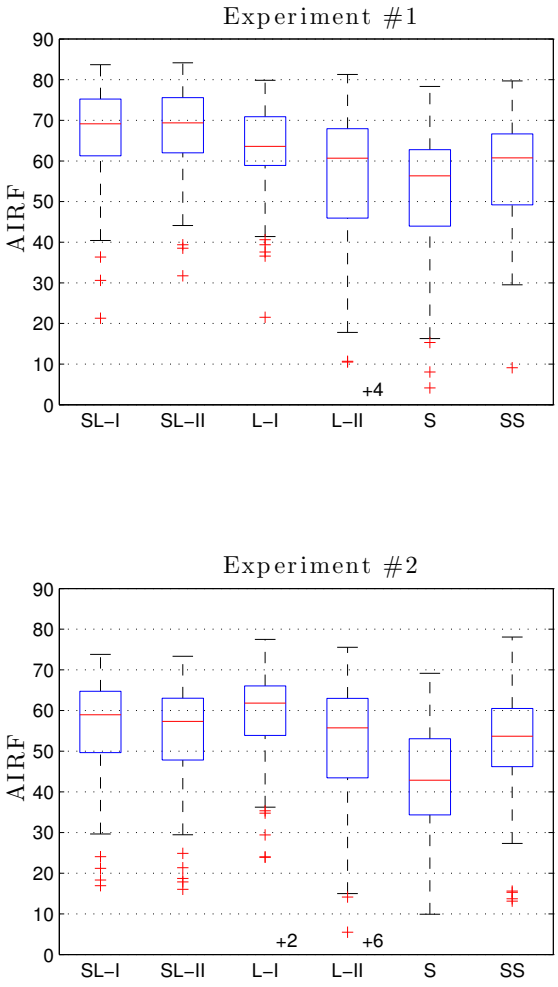

Fig. 4. Average impulse response fit (AIRF) obtained by the 6 estimators described in Section 6.

hence $P \leq \breve{P}$. Let $\varepsilon > 1$ and consider the decomposition

$$\breve{P} = D\bar{P}D \tag{56}$$

where $D \in \mathcal{S}_2(\mathbb{N}^+)$ is an infinite diagonal matrix with $\sqrt{t^{\alpha+\varepsilon}}e^{-\frac{\beta}{2}t}$ as t-th element in the main diagonal, and

$$[\bar{P}]_{ts} = \begin{cases} \kappa t^{-\varepsilon} & t = s \\ \star & \text{otherwise.} \end{cases} \tag{57}$$

Note that, $\bar{P}$ is strictly positive definite by construction. Moreover,

$$\operatorname{tr}(\bar{P}) = \kappa \sum_{t=1}^{\infty} t^{-\varepsilon} < \infty \tag{58}$$

because $\varepsilon > 1$. This means the sum of all the (nonnegative) eigenvalues of $\bar{P}$ is bounded, therefore its maximum eigenvalue, say λ, is bounded. Therefore, by (18) we have

$$\begin{aligned} c &\geq \mathbb{E}[\phi^\top P^{-1}\phi] \geq \mathbb{E}[\phi^\top \breve{P}^{-1}\phi] \\ &\geq \mathbb{E}[\phi^\top D^{-1}\bar{P}^{-1}D^{-1}\phi] \\ &\geq \lambda^{-1}\mathbb{E}[\phi^\top D^{-2}\phi] = \lambda^{-1}\sum_{t=1}^{\infty} t^{-(\alpha+\varepsilon)}e^{\beta t}[K]_{tt} \end{aligned} \tag{59}$$

accordingly $\sum_{t=1}^{\infty} t^{-(\alpha+\varepsilon)}e^{\beta t}[K]_{tt}$ is bounded and thus $[K]_{tt} \leq \bar{\kappa}_\varepsilon t^{\alpha+\varepsilon}e^{-\beta t}$, $t \in \mathbb{N}^+$, for some $\bar{\kappa}_\varepsilon > 0$. □

Proof of Theorem 4

We characterize the optimal solution by exploiting the duality theory. We consider the Lagrange function

$$\begin{aligned} \mathcal{L}(\mathbf{p}_\mathcal{I}, \tilde{\Gamma}, \lambda, \tilde{\Lambda}) =& \mathbf{H}(\mathbf{p}_\mathcal{I}) \\ &+ \frac{1}{2}\sum_{i,j=1}^{p} \tilde{\gamma}_{(i-1)p+j}(c_{ij} - \mathbb{E}[\|\check{s}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_s}}]) \\ &+ \frac{1}{2}\lambda(r - \sum_{i,j=1}^{p} \mathbb{E}[\|\check{l}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_l}}]) \\ &+ \frac{1}{2}\operatorname{tr}(\tilde{\Lambda}(Q - \mathbb{E}[\check{A}_l\check{A}_l^\top])) \end{aligned} \tag{60}$$

where $\tilde{\Gamma} = \operatorname{diag}(\tilde{\gamma}_1 \dots \gamma_{p^2})$ with $\tilde{\gamma}_i \geq 0$, $i = 1 \dots p^2$, $\lambda \geq 0$ and $\tilde{\Lambda} \in \overline{\mathcal{M}}^p_+$ are the Lagrange multipliers. It is not difficult to see that $\mathcal{L}$ is strictly concave over $\mathcal{P}$. Moreover, its unique maximum point is given by annihilating its first derivative. Therefore, we obtain

$$\begin{aligned} \mathbf{p}_\mathcal{I}(\check{\theta}_s, \check{\theta}_l) = \frac{1}{c}\exp\Bigg(&-\frac{1}{2}\sum_{i,j=1}^{p} \tilde{\gamma}_{(i-1)p+j}\|\check{s}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_s}} \\ &-\frac{1}{2}\lambda\sum_{i,j=1}^{p}\|\check{l}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_l}} - \frac{1}{2}\operatorname{tr}(\tilde{\Lambda}\check{A}_l\check{A}_l^\top)\Bigg) \end{aligned} \tag{61}$$

where c is the normalization constant. Let e_i, $i = 1 \dots p^2$, denote the j-th vector of the canonical basis of $\mathbb{R}^{p^2}$.

Then, we have

$$\sum_{i,j=1}^{p} \tilde{\gamma}_{(i-1)p+j} \|\check{s}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_s}} = \sum_{i,j=1}^{p} \tilde{\gamma}_{(i-1)p+j} \check{\theta}_s^\top (e_{(i-1)p+j} e^\top_{(i-1)p+j} \otimes P^{-1}_{\mathcal{I}_s}) \check{\theta}_s = \check{\theta}_s^\top (\tilde{\Gamma} \otimes P^{-1}_{\mathcal{I}_s}) \check{\theta}_s. \quad (62)$$

In similar way,

$$\sum_{i,j=1}^{p} \|\check{l}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_l}} = \check{\theta}_l^\top (I_{p^2} \otimes P^{-1}_{\mathcal{I}_l}) \check{\theta}_l. \quad (63)$$

Moreover, it is not difficult to see that

$$\mathrm{tr}(\tilde{\Lambda} \check{A}_l \check{A}_l^\top) = \check{\theta}_l^\top (\tilde{\Lambda} \otimes I_{p|\mathcal{I}_l|}) \check{\theta}_l. \quad (64)$$

Therefore, the optimal solution is such that $\mathbf{p}_{\mathcal{I}} = \mathbf{p}_{\mathcal{I}_s} \mathbf{p}_{\mathcal{I}_l}$, where

$$\mathbf{p}_{\mathcal{I}_s}(\check{\theta}_s) = \frac{1}{c_s} \exp\left(-\frac{1}{2} \check{\theta}_s^\top (\tilde{\Gamma} \otimes P^{-1}_{\mathcal{I}_s}) \check{\theta}_s \right)$$
$$\mathbf{p}_{\mathcal{I}_l}(\check{\theta}_l) = \frac{1}{c_l} \exp\left(-\frac{1}{2} \check{\theta}_l^\top (\lambda I_{p^2} \otimes P^{-1}_{\mathcal{I}_l} + \tilde{\Lambda} \otimes I_{p|\mathcal{I}_l|}) \check{\theta}_l \right)$$

with c_s and c_l normalization constants. Note that, $\mathbf{p}_{\mathcal{I}_s}$ and $\mathbf{p}_{\mathcal{I}_l}$ denote the marginal probability density of $\check{\theta}_s$ and $\check{\theta}_l$, respectively, Therefore, if the dual problem admits solution, $\check{\theta}_s$ and $\check{\theta}_l$ are independent, Gaussian with zero mean and covariance matrix

$$\check{K}_S = \tilde{\Gamma}^{-1} \otimes P_{\mathcal{I}_s}$$
$$\check{K}_L = (\lambda I_{p^2} \otimes P^{-1}_{\mathcal{I}_l} + \tilde{\Lambda} \otimes I_{p|\mathcal{I}_l|})^{-1}. \quad (65)$$

Moreover,

$$\mathbf{H}(\mathbf{p}_{\mathcal{I}}) = \frac{1}{2} \log\det(\check{K}_S) + \frac{1}{2} \log\det(\check{K}_L) + \text{ const. term.}$$

Therefore, the dual problem is equivalent to minimize the function

$$J(\tilde{\Gamma}, \lambda, \tilde{\Lambda}) = -|\mathcal{I}_s| \log\det(\tilde{\Gamma}) + \mathrm{tr}(\tilde{\Gamma} C) + \lambda r + \mathrm{tr}(\tilde{\Lambda} Q) - \log\det(\lambda I_{p^2} \otimes P^{-1}_{\mathcal{I}_l} + \tilde{\Lambda} \otimes I_{p|\mathcal{I}_l|}) \quad (66)$$

where

$$C = \mathrm{diag}(c_{11}, \ldots, c_{1p}, \ldots, c_{p1}, \ldots, c_{pp}). \quad (67)$$

Since $C \in \mathcal{M}^{p^2}_+$, $r > 0$ and $Q \in \mathcal{M}^p_+$, it is not difficult to see that J is lower bounded. It takes infinite value if and only if $\tilde{\Gamma}$ and/or λ and/or $\tilde{\Lambda}$ are not bounded (the formal proof follows the one of Proposition 5.1 in Ferrante et al. (2012)). Moreover, if $\tilde{\Gamma}$ and/or $\lambda I_{p^2} \otimes P^{-1}_{|\mathcal{I}_l|} + \tilde{\Lambda} \otimes I_{p|\mathcal{I}_l|}$ tend to be singular then J approaches infinity. Accordingly, we can restrict the search of the Lagrange multipliers over the closed and bounded set

$$\{(\tilde{\Gamma}, \lambda, \tilde{\Lambda}) \text{ s.t. } \varepsilon_1 I_{p^2} \le \tilde{\Gamma} \le M_1 I_{p^2}, \quad \varepsilon_2 \le \lambda \le M_2, \ \ \varepsilon_3 I_p \le \tilde{\Lambda} \le M_3 I_p\} \quad (68)$$

for some and $\varepsilon_1, M_1 > 0$ and $\varepsilon_2, \varepsilon_3, M_2, M_3 \ge 0$. The latter is a compact set because we are in a finite dimensional space. Since J is continuous over that set, by the Weiestrass Theorem, the dual problem admits solution. Accordingly, the kernel matrices for $\check{\theta}_s$ and $\check{\theta}_l$ solution to (28) does exist and are unique. Finally, setting $\gamma_i = \tilde{\gamma}_i^{-1}$ we obtain matrix K_S in the statement. □

Proof of Proposition 5

Consider the maximum entropy problem (28) where c_{ij} and Q have been replaced by $c^{(k)}_{ij}$ and $Q^{(k)}$ arbitrarily extracted from some sequences $\{c^{(k)}_{ij}\}_{k\ge 0}$, $c^{(k)}_{ij} > 0$, and $\{Q^{(k)}\}_{k\ge 0}$, $Q^{(k)} \in \mathcal{M}^p_+$, respectively, and such that $c^{(k)}_{ij} \to 0 \ \forall (i,j) \in \mathbf{C}$ and $Q^{(k)} v \to 0 \ \forall v \in \mathbf{Q}$ as $k \to \infty$. Let $[\check{\theta}^{(k)\top}_s \ \check{\theta}^{(k)\top}_l]^\top$ be the random vector solution to this maximum entropy problem. The corresponding kernel matrices are (29) and (30) where γ_i, λ and Λ have been replaced with $\gamma^{(k)}_i > 0$, $\lambda^{(k)} > 0$ and $\Lambda^{(k)} \in \mathcal{M}^p_+$ with $k \ge 0$. Then, substituting this random vector into (25) we obtain

$$c^{(k)}_{ij} \ge \mathbb{E}[\|\check{s}^{[ij]}\|^2_{P^{-1}_{\mathcal{I}_l}}] = \mathrm{tr}(\mathbb{E}[\check{s}^{[ij]} \check{s}^{[ij]^\top}] P^{-1}_{\mathcal{I}_s}) = |\mathcal{I}_s| \gamma^{(k)}_{(i-1)p+j} > 0. \quad (69)$$

Accordingly, if $(i,j) \in \mathbf{C}$ then $\gamma^{(k)}_{(i-1)p+j} \to 0$ as $k \to \infty$. Let $\check{A}^{(k)}_l$ be the matrix built from the maximum entropy random vector. Then, substituting it in constraint (27) pre- and post-multiplied by an arbitrary $v \in \mathbb{R}^p$, we have

$$\begin{aligned} v^\top Q^{(k)} v &\ge v^\top \mathbb{E}[\check{A}^{(k)}_l \check{A}^{(k)\top}_l] v = \mathbb{E}[\mathrm{tr}(v v^\top \check{A}^{(k)}_l \check{A}^{(k)\top}_l)] \\ &= \mathbb{E}[\check{\theta}^{(k)\top}_l (v v^\top \otimes I_{p|\mathcal{I}_l|}) \check{\theta}^{(k)}_l] \\ &= \mathrm{tr}((v v^\top \otimes I_{p|\mathcal{I}_l|}) \mathbb{E}[\check{\theta}^{(k)}_l \check{\theta}^{(k)\top}_l]) \\ &= \mathrm{tr}((v v^\top \otimes I_{p|\mathcal{I}_l|})(\lambda^{(k)} I_{p^2} \otimes P^{-1}_{\mathcal{I}_l} + \Lambda^{(k)^{-1}} \otimes I_{p|\mathcal{I}_l|})^{-1}) \\ &\ge v^\top (\lambda^{(k)} \mu I_p + \Lambda^{(k)^{-1}})^{-1} v p |\mathcal{I}_l| \\ &= v^\top (\Lambda^{(k)} - \Lambda^{(k)} (\Lambda^{(k)} + \frac{1}{\lambda^{(k)} \mu} I_p)^{-1} \Lambda^{(k)}) v p |\mathcal{I}_l| > 0 \end{aligned} \quad (70)$$

where we exploited (64) and $\mu > 0$ is the maximum eigenvalue of $P_{\mathcal{I}_l}$. Accordingly, if $v \in \mathbf{Q}$ then $\Lambda^{(k)} v \to 0$ as $k \to \infty$. □

Proof of Proposition 8

We use the same decomposition for $\breve{P}$ exploited in the proof of Proposition 2. Thus, we have

$$\begin{aligned} Q &\geq \mathbb{E}[A_l(P^{-1} \otimes I_p)A_l^\top] \geq \mathbb{E}[A_l(\breve{P}^{-1} \otimes I_p)A_l^\top] \\ &\geq \mathbb{E}[A_l(D^{-1} \otimes I_p)(\bar{P}^{-1} \otimes I_p)(D^{-1} \otimes I_p)A_l^\top] \\ &\geq \lambda^{-1}\mathbb{E}[A_l(D^{-2} \otimes I_p)A_l^\top] \\ &\geq \lambda^{-1}\sum_{t=1}^{\infty} t^{-(\alpha+\varepsilon)} e^{\beta t}\mathbb{E}[L_t L_t^\top] \end{aligned} \tag{71}$$

where λ is the maximum eigenvalue of $\bar{P}$ which is bounded. Hence, condition (36) implies that

$$\sum_{t=1}^{\infty} t^{-(\alpha+\varepsilon)} e^{\beta t}\mathbb{E}[|[L_t]_{ij}|^2] \tag{72}$$

is bounded for $i, j = 1 \dots p$. Accordingly, $\mathbb{E}[|[L_t]_{ij}|^2] \leq \bar{\kappa}_\varepsilon t^{\alpha+\varepsilon} e^{-\beta t}$, $t \in \mathbb{N}^+$, for some $\bar{\kappa}_\varepsilon > 0$. □

Proof of Theorem 9

The statement is proved by using the duality theory as in the proof of Theorem 4. In particular the Lagrange function is

$$\begin{aligned} \mathcal{L}(\mathbf{p}, \tilde{\Gamma}, \tilde{\Lambda}) =& \mathbf{H}(\mathbf{p}_{\mathcal{I}}) \\ &+ \frac{1}{2}\sum_{i,j=1}^{p} \tilde{\gamma}_{(i-1)p+j}(c_{ij} - \mathbb{E}[\|\breve{s}^{[ij]}\|^2_{P_{\mathcal{I}_s}^{-1}}]) \\ &+ \frac{1}{2}\operatorname{tr}(\tilde{\Lambda}(Q - \mathbb{E}[\breve{A}_l(P_{\mathcal{I}_l}^{-1} \otimes I_p)\breve{A}_l^\top])) \end{aligned} \tag{73}$$

where $\tilde{\Gamma} = \operatorname{diag}(\tilde{\gamma}_1 \dots \gamma_{p^2})$ with $\tilde{\gamma}_i \geq 0$, $i = 1 \dots p^2$, and $\tilde{\Lambda} \in \overline{\mathcal{M}}_p^+$. Then, it is not difficult to see that $\mathbf{p}_{\mathcal{I}}$ maximizing $\mathcal{L}$ is such that $\mathbf{p}_{\mathcal{I}} = \mathbf{p}_{\mathcal{I}_s}\mathbf{p}_{\mathcal{I}_l}$ where

$$\begin{aligned} \mathbf{p}_s &= \frac{1}{c_s}\exp\left(-\frac{1}{2}\breve{\theta}_s^\top(\tilde{\Gamma} \otimes P_{\mathcal{I}_s}^{-1})\breve{\theta}_s\right) \\ \mathbf{p}_l &= \frac{1}{c_l}\exp\left(-\frac{1}{2}\breve{\theta}_l^\top(\tilde{\Lambda} \otimes I_p \otimes P_{\mathcal{I}_l}^{-1})\breve{\theta}_l\right) \end{aligned} \tag{74}$$

with c_s and c_l normalization constants. Therefore, if the dual problem admits solution, $\breve{\theta}_s$ and $\breve{\theta}_l$ are independent, Gaussian, with zero mean and covariance matrix, respectively,

$$\begin{aligned} K_S &= \tilde{\Gamma}^{-1} \otimes P_{\mathcal{I}_s} \\ K_L &= \tilde{\Lambda}^{-1} \otimes I_p \otimes P_{\mathcal{I}_l}. \end{aligned} \tag{75}$$

Finally, the dual problem is equivalent to minimize the function

$$\begin{aligned} J(\tilde{\Gamma}, \tilde{\Lambda}) = &-|\mathcal{I}_s| \log\det(\tilde{\Gamma}) + \operatorname{tr}(\tilde{\Gamma}C) \\ &- p|\mathcal{I}_l| \log\det(\tilde{\Lambda}) + \operatorname{tr}(\tilde{\Lambda}Q) \end{aligned} \tag{76}$$

where C has been defined in (67). In this case the search of the minimum can be restricted to the compact set

$$\{(\tilde{\Gamma}, \tilde{\Lambda}) \text{ s.t. } 0 < \tilde{\Gamma} \leq M_1 I_{p^2},\ \ 0 < \tilde{\Lambda} \leq M_2 I_p\} \tag{77}$$

for some $M_1, M_2 > 0$. Moreover, J is continuous over this set. Thus, by the Weiestrass Theorem the dual problem admits solution. Finally, setting $\gamma_i = \tilde{\gamma}_i^{-1}$ and $\Lambda = \tilde{\Lambda}^{-1}$ we obtain the kernel matrices in the statement. □

Proof of Proposition 10

The proof follows the same lines of the one of Proposition 5.

Proof of Proposition 13

It is sufficient to observe that Problem (45) can be rewritten as

$$\hat{\theta} = \arg\min_{\tilde{\theta} \in \mathcal{H}_{\tilde{K}}} \|\mathrm{y}^+ - \Phi\tilde{\theta}\|^2_{\tilde{\Sigma}^{-1} \otimes I_N} + \|\tilde{\theta}\|^2_{\tilde{K}^{-1}} \tag{78}$$

where $\tilde{\theta} = \begin{bmatrix} \theta_s^\top & \theta_l^\top \end{bmatrix}^\top$, the regression matrix is $\tilde{\Phi} = \begin{bmatrix} \Phi & \Phi \end{bmatrix}$, and $\mathcal{H}_{\tilde{K}}$ is the reproducing Hilbert space of deterministic functions on $\mathbb{N}^+$ with kernel

$$\tilde{K} = \begin{bmatrix} K_S & 0 \\ 0 & K_L \end{bmatrix}. \tag{79}$$

Then, the statement follows from Proposition 3 in Pillonetto et al. (2011). More precisely, it is not difficult to see that Proposition 3 still holds when the covariance matrix of the noise vector is not diagonal. □

References

Øksendal, B. (1998). *Stochastic differential equations: an introduction with applications*. Universitext. Berlin, Heidelberg, New York: Springer.

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*, 716–723.

Aravkin, A., Burke, J., Chiuso, A., & Pillonetto, G. (2014). Convex vs non-convex estimators for regression and sparse estimation: the mean squared error properties of ard and glasso. *Journal of Machine Learning Research*, *15*, 217–252.

Aronszajn, N. (1950). Theory of reproducing kernels. *Trans. Amer. Math. Soc.*, *68*, 337–404.

Bonettini, S., Chiuso, A., & Prato, M. (2015). A scaled gradient projection method for Bayesian learning in dynamical systems. *SIAM Journal on Scientific Computing (accepted)*, .

Carli, F., Chen, T., & Ljung, L. (2014). Maximum entropy kernels of system identification. *ArXiv*, .

Chandrasekaran, V., Parrilo, P., & Willsky, A. (2010). Latent variable graphical model selection via convex optimization. *Annals of Statistics (with discussion)*, *40*, 1935–2013.

Chen, T., Ardeshiri, T., Carli, F. P., Chiuso, A., Ljung, L., & Pillonetto, G. (2015). Maximum entropy properties of discrete-time first-order stable spline kernel. *CoRR*, *abs/1504.03253*.

Chen, T., Ohlsson, H., & Ljung, L. (2012). On the estimation of transfer functions, regularizations and gaussian processes-revisited. *Automatica*, *48*, 1525–1535.

Chiuso, A., & Pillonetto, G. (2012). A Bayesian approach to sparse dynamic network identification. *Automatica*, *48*, 1553–1565.

Choi, M., Chandrasekaran, V., & Willsky, A. (2010). Gaussian multiresolution models: Exploiting sparse markov and covariance structure. *IEEE Transactions on Signal Processing*, *58*, 1012–1024.

Cover, T., & Thomas, J. (1991). *Information Theory*. New York: Wiley.

Deistler, M., Scherrer, W., & Anderson, B. (2015). The structure of generalized linear dynamic factor models. In *Empirical Economic and Financial Research* (pp. 379–400). Springer.

Deistler, M., & Zinner, C. (2007). Modelling high-dimensional time series by generalized linear dynamic factor models: An introductory survey. *Communications in Information & Systems*, *7*, 153–166.

Ding, M., Chen, Y., & Bressler, S. (2006). Granger causality: basic theory and application to neuroscience. *Handbook of time series analysis: recent theoretical developments and applications*, (p. 437).

Ferrante, A., Pavon, M., & Zorzi, M. (2012). A maximum entropy enhancement for a family of high-resolution spectral estimators. *IEEE Transaction on Automatic Control*, *57*, 318–329.

Goodwin, G., Gevers, M., & Ninness, B. (1992). Quantifying the error in estimated transfer functions with application to model order selection. *IEEE Transactions on Automatic Control*, *37*, 913–928.

Granger, C. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, *37*, 424–438.

Heij, C., Scherrer, W., & Deistler, M. (1997). System identification by dynamic factor models. *SIAM Journal on Control and Optimization*, *35*, 1924–1951.

Jorgesen, P., Kornelson, K., & Shuman, K. (2011). *Iterated Function Systems, Moments, and Transformations of Infinite Matrices*. American Mathematical Society.

Lauritzen, S. (1996). *Graphical Models*. Oxford: Oxford University Press.

Liegegois, R., Mishra, B., Zorzi, M., & Sepulchre, R. (2015). Sparse plus low-rank identification in neuroimaging times series. In *Proceedings of the IEEE Conference on Decision and Control (accepted)*. Osaka.

Ljung, L. (1999). *System Identification: Theory for the User*. New Jersey: Prentice Hall.

Nicolao, G. D., Ferrari-Trecate, G., & Lecchini, A. (1998). MAXENT priors for stochastic filtering problems. In *Mathematical Theory of Networks and Systems*. Padova, Italy.

Pillonetto, G., Chiuso, A., & De Nicolao, G. (2011). Prediction error identification of linear systems: A nonparametric gaussian regression approach. *Automatica*, *47*, 291–305.

Pillonetto, G., & De Nicolao, G. (2010). A new kernel-based approach for linear system identification. *Automatica*, *46*, 81–93.

Pillonetto, G., Dinuzzo, F., Chen, T., De Nicolao, G., & Ljung, L. (2014). Kernel methods in system identification, machine learning and function estimation: A survey. *Automatica*, *50*, 657–682.

Poggio, T., & Girosi, F. (1990). Networks for approximation and learning. *Proceedings of the IEEE*, *78*, 1481–1497.

Prando, G., Pillonetto, G., & Chiuso, A. (2015). Maximum entropy vector kernels for mimo system identification. *arXiv:1508.02865*, .

Rasmussen, C., & Williams, C. (2006). *Gaussian Processes for Machine Learning*. The MIT Press.

Schwarz, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics*, *6*, 461–464.

Söderström, T., & Stoica, P. (1989). *System Identification*. New Jersey: Prentice Hall.

Wipf, D. (2012). Non-convex rank minimization via an empirical Bayesian approach. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence* (pp. 914–923).

Wipf, D., & Nagarajan, S. (2010). Iterative reweighted ℓ_1 and ℓ_2 methods for finding sparse solutions. *J. Sel. Topics Signal Processing*, *4*, 317–329.

Zorzi, M., & Sepulchre, R. (2015). AR identification of latent-variable graphical models. IEEE Transactions on Automatic Control (to appear).